Chapter 3 Lane Transition

3.1 Introduction

It is obvious to think about lane transition only in terms of moving a vehicle from one driving lane to another adjacent lane. While this task is the most evident, there are many others that are similar in nature which require a subset of or minor extension to the capabilities required for lane transition. In addition to presenting three techniques for lane transition, this chapter introduces algorithms for exit and entrance lane detection and traversal, short distance obstacle avoidance maneuvering, and extended driving outside of the typical lane boundaries.

All of these tasks are assuming increasing importance because of the requirements for new advanced traffic control projects like the Automated Highway System (AHS) in the United States and PROMOTE in Europe. For the AHS, moving vehicles into and out of automated lanes and maneuvering to avoid disabled or non-automated vehicles or other obstacles are the primary situations where lane transition and related maneuvers are needed.

Technically, these tasks are important because they represent tactical level driving tasks which have not been studied as closely as low level tasks like lane keeping and obstacle detection. In order to be successfully accomplished, all of the tasks require a sequence of actions to be taken - a plan. This type of behavior is difficult to coax from low level systems, which are characterized by fast response to features in their surroundings. Yet traditional planning is often rigid and difficult to integrate into autonomous systems which function in the real world. The algorithms presented in this chapter circumvent the problem by framing the task in a way that preserves the flexibility and robustness of the low level lane keeping system, yet exhibits enough control to permit successful execution.

This behavior is implemented using active placement and control of virtual cameras, intelligent interpretation of the lane keeping system's response to the images created from the virtual cameras, and a simple road model. Using the information contained in the model, the virtual camera views [22] can be positioned which allow existing lane keeping systems to be used for these maneuvers. The techniques presented are based on controlling the lateral placement of virtual camera views using the road model and control scheme for guidance. They are not strictly tied to any specific lane keeping system, only requiring that the system take images as input and produce a point on the road to drive over and a measure of its internal confidence in this point as output. When implemented using a slightly modified version of the ALVINN [40] lane keeping system, the techniques have been able to autonomously navigate our testbed vehicle, the Navlab 5 [24], between lanes of a rural interstate highway, onto exit ramps and around obstacles, as well as to offset the vehicle within the current driving lane.

3.2 Road Model

All of the techniques described in this chapter require a very simple lane model. The model that the techniques use is that the lane centers have constant separation. This implies that the lanes are parallel and have a constant width. Because this model is so simple, it can either be pre-specified at start up or can be passed to the system from some other, higher level module.

3.3 Other Work

There has been a significant amount of research published describing how people change lanes as well as identifying theoretically optimal control strategies which could or should be used to autonomously control a vehicle in a lane change maneuver [15][17][18][19][36][49]. Unfortunately, many of the researchers didn't have the facilities or equipment to test their results outside of the lab.

An early and successful approach to lane transition is presented in [2]. In this approach, a vision-based model-driven lane keeping system is used to find lane markings [8]. When a lane change is initiated, instead of driving exactly between the left and right lane markings, the system offsets the vehicle in the direction of the destination lane. The source lane markings are tracked until either they leave the camera field of view or when the vehicle crosses the center line of the road. When either of these occur, the system looks for the destination lane markings, using a road model like the one mentioned in Section 3.2. Once it acquires the destination lane markings, the system recenters the vehicle in this new lane. This technique is very similar to the Incremental View Lane Transition algorithm, which will be presented in Section 3.4.4, except that it is more closely tied to their lane keeping system.

Another system which executes lane change maneuvers is ROBIN. (See Section 5.2 for a detailed description of ROBIN.) Although ROBIN uses virtual cameras to enhance its driving capabilities, it does not use any of the active virtual camera control techniques presented later in this chapter. The method that it uses is very similar to the network switching method described in Section 3.4.2. The ROBIN system uses two radial basis function networks which are trained to drive in the right and left lanes of a two lane interstate highway. When the system is transitioning lanes, a slide behavior is activated. When this behavior is activated no network is controlling the vehicle and it is essentially driving blind between the lanes. When the slide is complete the destination lane network is activated. The lateral slide distance is set so that when the destination lane network is activated after the slide is complete, it will be able to recognize the situation and finish reentering the vehicle in the new lane. If the lateral slide distance is too short, the destination lane network becomes confused or acts abruptly when it was activated. This method works well at low speeds on straight sections of highway, but does not perform adequately on curves. It has not been tested at typical highway speeds [45].

In a related area, experimental work on lane structure detection was reported in [56]. The system uses video images to detect the presence of additional lanes. This system was not used to control the vehicle, but rather to provide geometric information about the environment around the vehicle.

3.4 Lane Change

Many systems have been created which can keep a vehicle within a driving lane, but little experimental work has been reported that describes methods to transition an autonomous vehicle between lanes. This task is assuming increasing importance because of the requirements for new advanced traffic control projects like the Automated Highway System (AHS). For the AHS, moving vehicles into and out of automated lanes is the primary situation that lane change maneuvers are needed and one which is on the critical path for deployment of full scale AHS.

Three techniques to accomplish lane transition have been investigated. The most basic involves intelligently switching between two trained ALVINN networks. The other techniques use active control of virtual camera views to move the vehicle into the destination lane.

3.4.1 Control Scheme

All of the lane transition techniques presented in this chapter are implemented using a linear control scheme to determine the appropriate way to move to the next step of the lane transition maneuver. This scheme was chosen because it provides a simple, yet effective, control mechanism for autonomous lane transitions using the Navlab 5. The techniques presented in this chapter are not constrained to using a linear scheme, but because it leads to reliable, intuitive lane transitions when used with the ALVINN lane-keeping system, other schemes were not investigated.

When using this scheme with the Incremental Network Switching technique it defines how to integrate two network output displacements into a single displacement which can be used to compute a transition point to drive over. Briefly, if a transition of this type has 10 steps, during the first step 90% of the output displacement of the driving lane network will be combined with 10% of the destination lane output displacement to compute a final displacement. On the second step, the percentages would be 80% and 20%. See Section 3.4.3 for a more detailed explanation of this technique. For the Incremental View and Dual View Lane Transition techniques, described in Section 3.4.4 and Section 3.4.5, the scheme defines how virtual camera views are to be moved at each step in the transition. For these techniques, the lane separation distance is divided by the desired number of lane transitions steps. The resultant step distance is then used to update the virtual camera views appropriately.

All of the lane transition techniques except those presented in Section 3.5, Lateral Offset Driving, require two trained networks. This requirement is not a function of the techniques, but rather of the underlying lane keeping system. Because ALVINN learns to key off features at particular locations in the input image, a single network can only drive in lanes or on roads whose feature are consistent with those learned. For the left and right lanes of an interstate highway, this is not the case, thus the need for two networks.

3.4.2 Lane Transition by Network Switching

Because ALVINN has the ability to recenter the vehicle in the driving lane even if it is displaced by a large distance, it is possible to transition from the current driving lane to the destination lane by simply switching to a network that was trained to drive in the destination lane. When presented with an image of the road while in the driving lane, this network would recognize that the vehicle was offset by a large amount from the proper driving position (in the destination lane) and would produce a sharp turn to recenter the vehicle. This type of lane transition is inadequate for two reasons. First, it is a very unsafe maneuver. At 25 meters/second, the lateral acceleration resulting from the sharpest turn ALVINN typically produces for highway driving is approximately 4.0 m/s^2. Even though the system cycles at 12-15 Hz, it may not recover quickly enough to prevent the vehicle from driving through the destination lane and off the other side of the road. Second, this type of maneuver would be uncomfortable and frightening to any passenger in the vehicle. A more controlled approach is desired in which the vehicle slowly transitions from one lane to the other, monitoring its performance along the way. To accomplish this goal, a modification of the network switching idea, along with two techniques based on controlled placement of virtual camera views, have been developed.

3.4.3 Incremental Network Switching

A method to transition lanes which is based on the simple network switching idea described above is called Incremental Network Switching (INS). Instead of immediately switching from the driving lane network to the destination lane network, this method combines the output displacements of the driving and destination lane networks to slowly transition the vehicle from the driving to the destination lane.

The INS method requires two trained networks, one for each lane. When a lane transition is initiated, road images are passed not only to the current driving lane network, but also to the network trained to drive in the destination lane. Each network responds by producing the displacement at the lookahead distance that would be required to recenter the vehicle in its respective lane.

Initially, the destination lane network's displacement is large, indicating that a hard turn is necessary. But instead of using this displacement alone, the driving lane network's displacement is also utilized to compute a target displacement, which will determine how to steer the vehicle. This is done by incrementally adjusting how much contribution the displacement associated with the driving and destination lane networks have to the target driving displacement.

Initially, the target displacement is very similar to the driving lane displacement. This prevents the sharp transition into the destination lane which is present in the basic lane switching method. But because a small part of the destination lane displacement is used, the vehicle does begin to move toward the destination lane. In subsequent iterations, the amount of the displacement contributed by the destination lane increases and the amount contributed by the original driving lane decreases. When the amount contributed by the destination lane becomes significant, its output displacement is smaller because the vehicle is much closer to the center of the destination lane. This means that sudden, sharp movements are again averted. As the transition ends, the destination lane output displacement is the major contributor to the target driving displacement. Although the original driving lane network is producing a displacement which would cause the vehicle to return sharply to the original driving lane, its contribution is small and has only a limited effect. For these experiments, the lane transition was specified to be accomplished in 30 increments. This means that at each step, the contribution of the driving lane displacement was reduced by 1/30 while the contribution of the destination lane displacement was increased by 1/30.

Two constraints were used to ensure that both networks were performing adequately during the lane transition maneuver. The IRRE confidence measure [40] produced by each network was the first constraint. This measure is an indicator which is strongly correlated to satisfactory driving behavior. In order for an INS lane transition to proceed, the IRRE measure produced by either

Left Lane Right Lane

Left Lane Right Lane

Figure 16 INS processing.

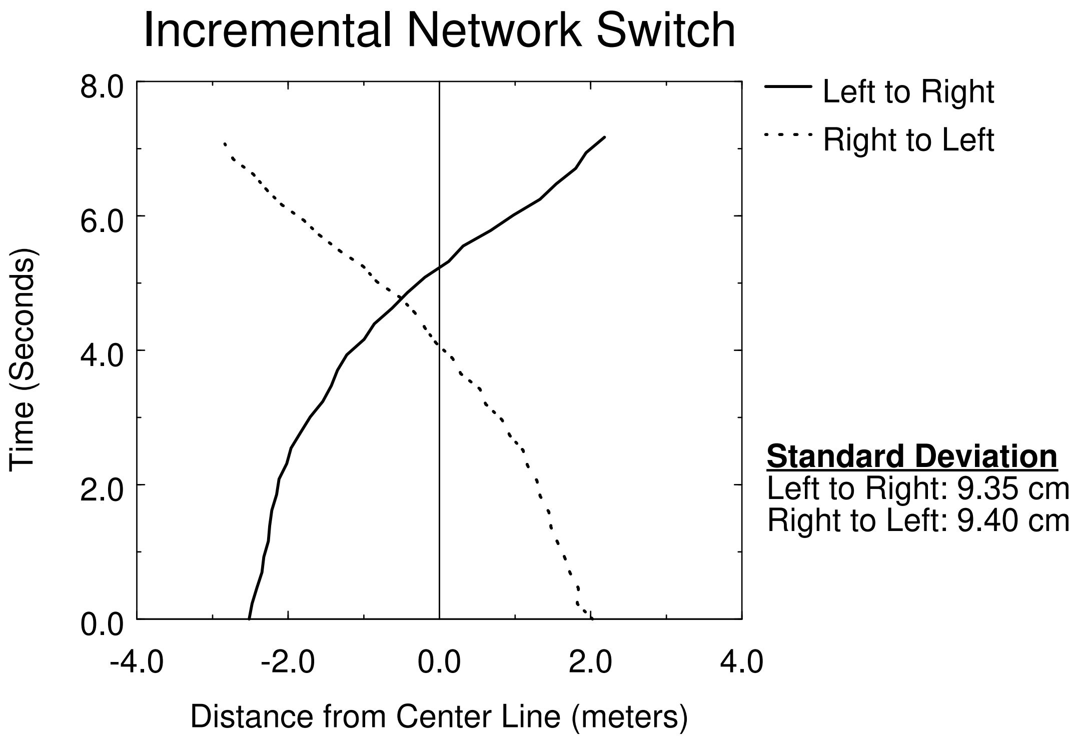

Figure 17 Average vehicle path during Incremental Network Switch Lane Transition.

of the networks was required to be greater than a threshold value. The reason only one network was required to produce a high IRRE value was the initially large lateral offset from the destination lane. This offset could be significant enough to cause the network trained for the destination lane to have difficulty correctly reconstructing the input image, but at the same time produce a reasonable steering response. Basically, it is possible that the network was never presented with images shifted similarly during training. Typically however, both networks produced satisfactory measures, with the measure increasing (or decreasing) as the vehicle moved into (or out of) the lane a network was trained to drive in.

The second constraint was related to the geometric relevance of the network's output displacements. In general, the difference between the displacements should be equal to the lane separation distance, specified in the road model. The driving lane network should specify the center of the driving lane while the destination lane output specifies the center of the destination lane. In order for this constraint to be satisfied, the difference between displacements was required to be within 40% of the lane separation distance. This figure allowed for inadequacy in the network's output to represent all displacements while still prohibiting obviously false combinations of displacements.

Figure 16 shows two images taken at the beginning and end of a left to right INS lane transition. In both images, the virtual camera view footprint on the ground plane is shown immediately in front of the vehicle. Items associated with each of the networks used for the INS are grouped in the lower left corner of each image. Each grouping contains a preprocessed ALVINN image created from the virtual camera view, ALVINN's output steering direction vector (shown above the preprocessed image), a bar graph representation of the IRRE confidence metric (shown to the right of the preprocessed image), and the network name (shown above the output steering direction vector). The grouping for the left lane network is on the left, while the one associated with the right lane is on the right. Also shown in each image as a grey box inside the virtual camera view footprint is the point on the road that the system is currently steering toward.

In the top image, the steering output of the left lane network - the one the vehicle is currently in - is almost straight ahead, while the output of the right lane network indicates a sharp right turn is required to recenter the vehicle. In the bottom image the situation is reversed. The output of the right lane network is nearly straight ahead while the output steering direction of the left lane network signifies that a hard left turn is necessary to recenter the vehicle.

Of the 38 lane transitions (19 in each direction) that were attempted, 37 succeeded. The results of these transitions are shown in Figure 17. This figure shows the average vehicle position with respect to the center line of the road versus time. (The average distance for completion of an INS lane transition was 156 meters, while the average vehicle speed was about 22.2 meter/second. Also note that the standard deviation for each transition type is the average deviation across all time slots within the transition.) The most obvious feature in this graph is that the two trajectories do not cross at the center line of the road. Initially it was suspected that this was due to the vehicle not being properly centered during training of the left lane network. But as will be shown later, this assumption proved to be false. The exact cause of the anomaly is not known, but it is suspected that it is related to some small difference in how the left and right lane networks respond to images which produce large displacements.

A somewhat subtler feature that is visible in the figure is the that the trajectories cross later in time than would be expected. If the transition was perfect, the trajectories would cross at about the 3.5 second mark, but using this method, they cross much later. This characteristic manifests itself

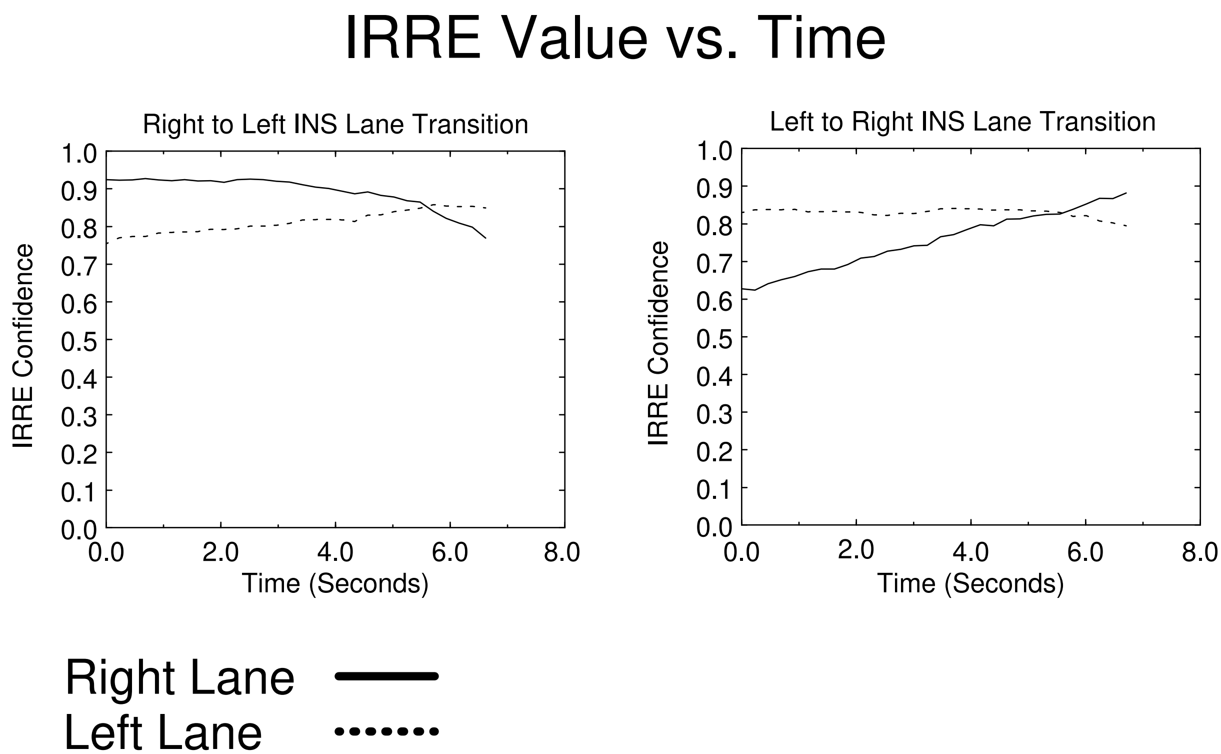

Figure 18 IRRE values during INS Lane Transitions.

during experiments by causing the vehicle to "dive" into the destination lane. Specifically, instead of moving smoothly into the destination lane, the vehicle's lateral velocity increases as the lane transition continues. This leads to overshoot in the destination lane after the maneuver has completed because of the somewhat long correction time required for smooth driving at high speeds. One explanation for this behavior relates to each network's inability to produce the large displacements that are required to locate the center of each lane during the beginning and end of the INS. Initially, the destination lane network produces the largest steering response it is capable of creating in the direction of the destination lane. Although this response is in the correct direction, it is not large enough to accurately locate the destination lane center. The result is that the vehicle initially moves slower toward the destination lane than would be anticipated. As the transition continues and the vehicle nears the center of the road, both networks begin to produce accurate steering responses and the vehicle moves laterally as expected. But once the vehicle has moved far enough from the original driving lane, the network associated with that lane cannot produce a large enough steering response to accurately locate the center of the lane. In this case, the target displacement is greater than what the geometry dictates, and the vehicle accelerates into the destination lane.

Figure 18 shows the average IRRE confidence values for the right and left lane ALVINN networks during right to left and left to right Incremental Network Switch lane transitions. Note that the confidence value of the destination lane network increase throughout the switch, while the original driving lane network's confidence decreases. Also, throughout the transition, the average IRRE confidence value was high for each network, even when the vehicle was not in the lane that the network was trained to drive in. This was a result of both road/shoulder boundaries being visible in the virtual camera image ALVINN used. Because the boundaries were partially visible in the image when the vehicle was in either lane, a high IRRE confidence value resulted. But as the graph shows, the network's ability to reconstruct the input image (which leads to high IRRE values) was beginning to diminish as the vehicle neared the end of the lane transition.

3.4.4 Incremental View Lane Transition

The Incremental View Lane Transition (IVLT) method uses two virtual camera views to move the vehicle from the driving lane to the destination lane. Unlike the INS method, the IVLT technique accomplishes lane transition by intelligently controlling the placement of these views and using ALVINN's response to the images created from them. The central idea behind this technique is that virtual camera views can be used as a tool to check if the upcoming driving situation can be correctly handled by ALVINN before actually encountering the situation. If it is determined that the situation can be handled, ALVINN's network output, produced from the images created by the virtual camera views, is used to control the vehicle in a manner which produces the desired lane transition.

Two virtual camera views are created for the IVLT technique. During normal operation before a lane transition has begun, only one virtual camera view, the Driving View, is used to keep the vehicle in the source driving lane. Initially, the Driving View is at the same location as the view which was used during training of the source lane network. When an IVLT is initiated, an additional view, called the Test View, is created. The only difference between the Driving View and the Test View is that the Test View is offset a small lateral distance, in the opposite direction of the desired lane change, from the Driving View. Lane transition is accomplished by switching the source of the images which ALVINN uses for vehicle control from the Driving View to the Test View. The lateral offset of the Test View leads to images in which the vehicle seems to be improperly aligned with the driving lane. Because ALVINN's neural network is trained to recover from these situations, it responds with steering commands that attempt to recenter the vehicle. Once the vehicle has been recentered in the virtual image, a new, shifted, Test View is created. As a side effect of the shifted views and resultant steering commands, the vehicle is incrementally moved through several lane transition locations and into the destination lane. The relative location of the Driving and Test Views for a right to left lane transition is shown in Figure 19.

The IVLT method requires two networks be trained, one network for driving in the source lane and another for the destination lane. The source lane network is used by the IVLT algorithm when the vehicle is moving from the center of the source lane to the boundary between lanes. When the vehicle reaches this point in the transition, the IVLT algorithm uses the destination lane network to finish the transition. Although the IVLT method requires two networks, only one is used to control the vehicle at any time during the transition. While the system is using the image created from the Driving View to servo the steering wheel, it uses the Test View image to determine if the system will be able to reliably move the vehicle to and continue control at the next transition location.

To determine if the next transition location is satisfactory, the Test View is used by the IVLT technique to compute both internal and external confidence metrics. These metrics can be used to determine if moving to the new lane location is advisable before actually executing the move. The IRRE confidence computed from ALVINN's output when presented with an image created from the Test View is used as an intrinsic measure of how well the network thinks it will be able to drive if the Test View is switched to the Driving View. Geometric knowledge about the lateral offset distance of the Test View from the Driving View, along with the output displacement of the Test and Driving Views, is used as the external metric. Because the system uses this geometric constraint in addition to the purely intrinsic IRRE measure, it is less susceptible to transitory features in the input image which can cause improper driving behavior yet not produce substantially lower IRRE values.

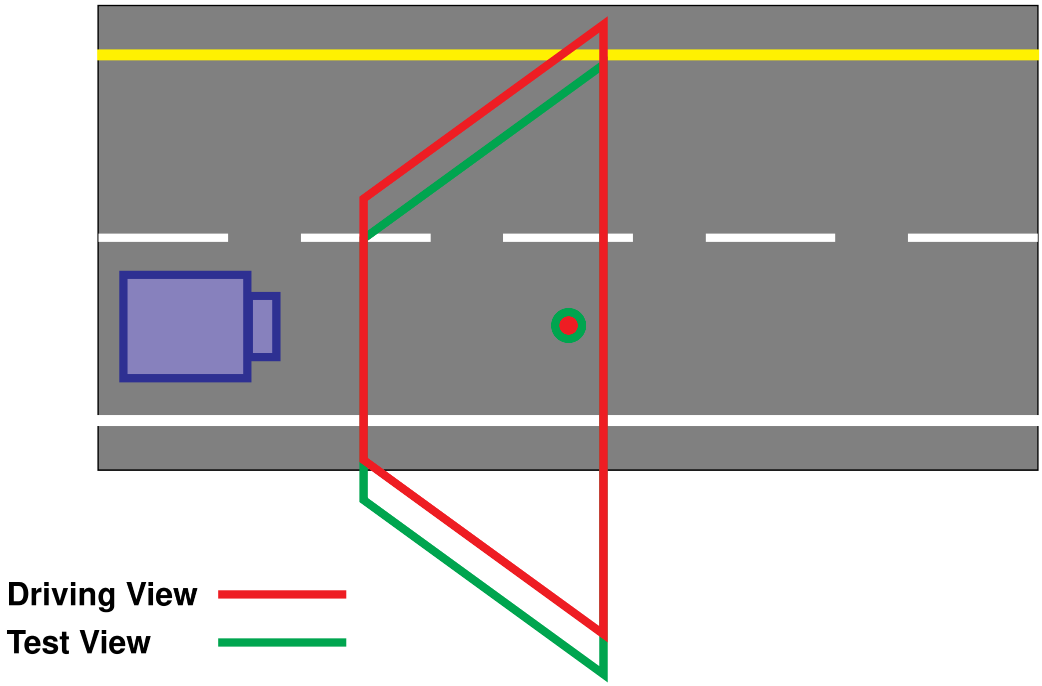

Figure 19 Driving and Test Views for a right to left lane transition.

Before the Test View can be created, the appropriate lateral offset distance is computed. The lateral offset distance applied to the Test View is determined by the lane model as well as the desired number of intermediate lane transition positions. Because the view associated with each lane transition position must be used for a fixed period of time - while the system is determining if moving to the next location is advisable - the selection of the number of lane transition locations is proportional to the transition rate of the vehicle. Typical lane separation distances of 3.2 to 4.0 meters, along with 10-20 intermediate lane transition positions, lead to controlled and comfortable IVLT's.

If the vehicle is to transition into a lane to its left, the Test View is offset to the right of the Driving View. The opposite holds for lane transitions to the right. An example of initiating a right to left lane transition in shown in Figure 19. The network associated with the Test View is in most cases the same as the one associated with the current Driving View. Because the Test View is laterally offset, the image it creates looks as if it was taken from a vehicle that was not centered in the driving lane. Because of this, when the Test View becomes the Driving View after the internal and external reliability constraints have been met, ALVINN will compensate by steering the vehi-

Driving View **Test View**

Figure 20 IVLT processing.

cle so that the road is centered in the image. The effect is to move the vehicle by the offset amount in the opposite direction of the offset.

Before actually moving from the current lane transition position, the IVLT method estimates ALVINN's ability to drive at the next location. It does this by monitoring the internal and external reliability metrics produced by ALVINN's neural network and the geometric constraints. To compute the external metric, the system computes the difference between the output displacement of the Test View network and the Driving View network. This difference is compared to the lateral offset distance of the Test View. If the difference between these figures is below some threshold (typically 40% of the Test View Offset), and the IRRE confidence is above another threshold (typically 0.40 out of 1.0) for a few (typically 2) Test View images, the system switches to using it as the Driving View and creates another Test View, again offset by the same distance from the new Driving View. This cycle repeats until the vehicle has reached the midway point of the lane transition. At this location, instead of using the same network as the Driving View, the Test View uses

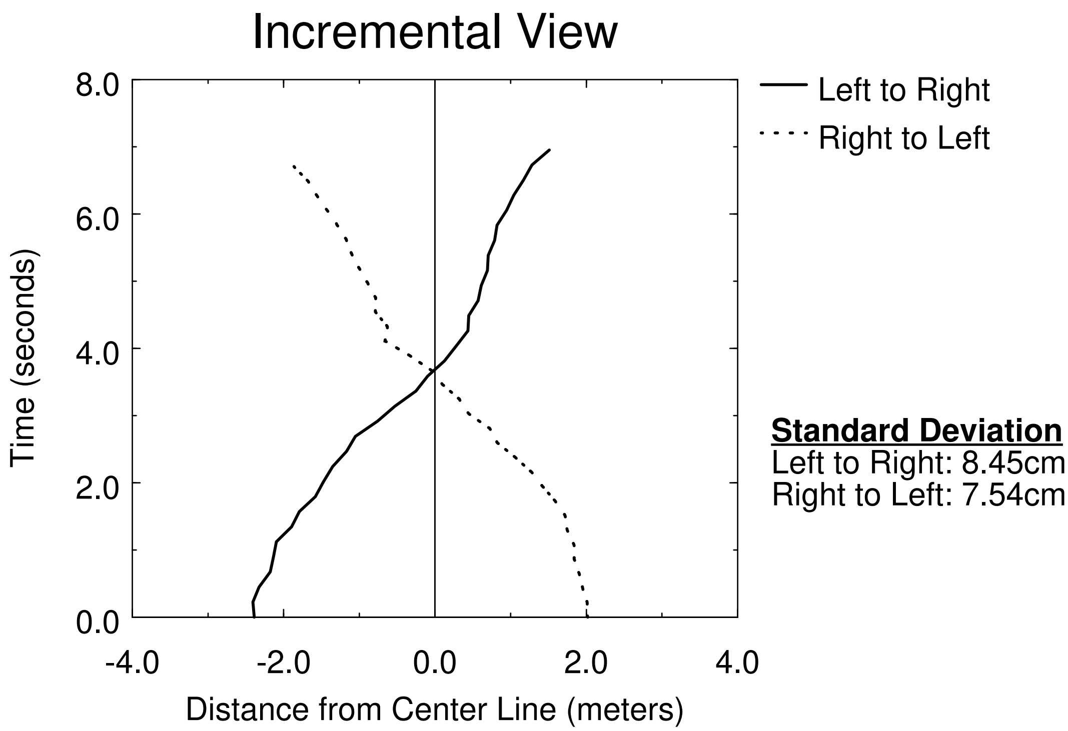

Figure 21 Average vehicle path during Incremental View Lane Transitions.

the network associated with the destination lane. Also at this step, the Test View offset distance from the Driving View is modified so that it correctly images the destination lane.

Figure 20 shows an image taken at the beginning of a left to right IVLT. The Driving View footprint is visible as the darker trapezoid, while the Test View footprint is displayed as the lighter trapezoid, offset a small distance to the left of the Driving View footprint. The preprocessed image, steering direction, and IRRE confidence groups are present in the lower left corner of the image, with the Driving View group positioned on the left.

ALVINN's steering response to the preprocessed image created from each view is nearly identical. But upon close examination, the steering output vector associated with the Test View is shifted slightly right of its position for the Driving View. This indicates that a a shallow right turn will be executed when the Test View is switched to the Driving View, leading to a left to right lane transition. Also note that the IRRE confidence associated with both preprocessed images is very high, indicating that ALVINN believes that it has encountered similar scenes during training.

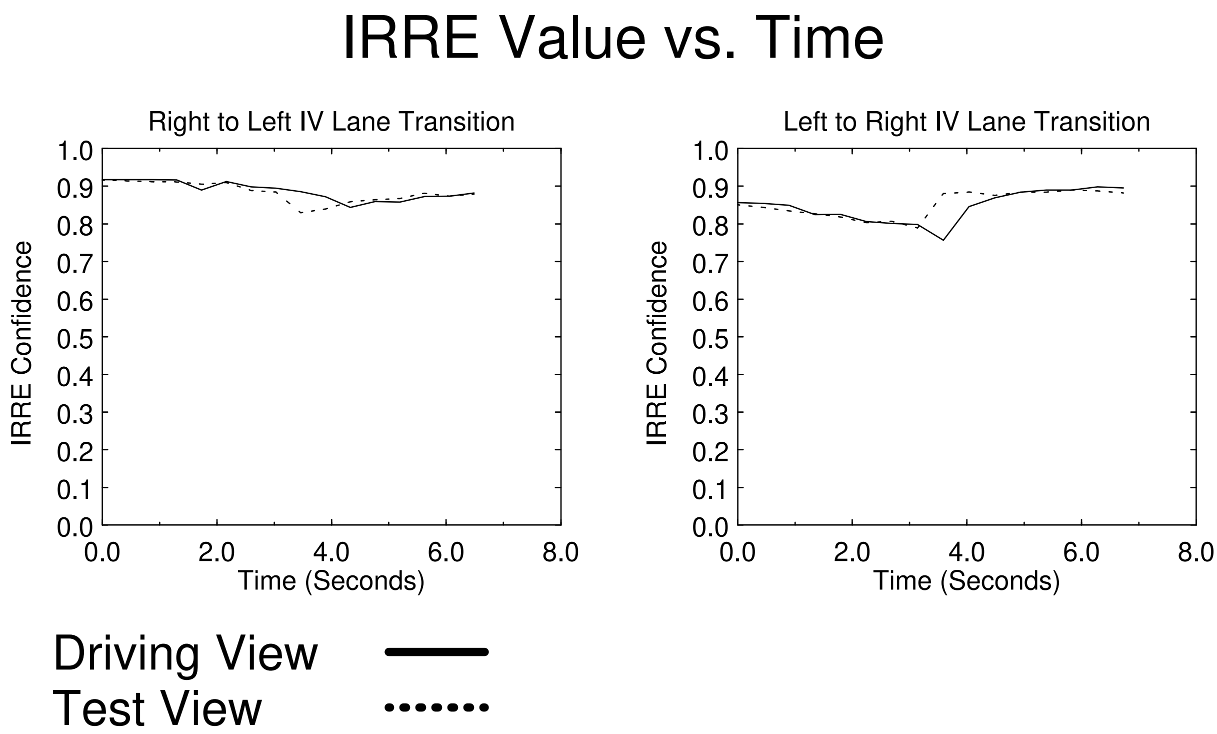

Figure 22 IRRE values during IVLT's.

All 32 IVLT's that were attempted to verify this lane transition method were successfully completed. Like the INS method, half were left lane to right lane transitions. The same two networks that were used for the INS method experiments were for these tests. The results of these transitions are shown in Figure 21. (The average distance for completion of an IVLT was 151 meters, while the average vehicle speed was about 21.6 meter/second.) There are three things to note about this graph. First, the average trajectory lines cross at the center line of the road. This is contrary to what occurred for the INS technique, which was thought to indicate improperly trained networks. Second, the trajectories are not as smooth as those in the INS method. This is a result of intermittent steering output due to the reduced cycle rate when a Test View is being processed. The last characteristic to note is that the IVLT method relinquishes control before the vehicle is at a typical driving position in the destination lane. This characteristic was unanticipated, but upon closer inspection it was determined to be the result of at least two contributing algorithmic factors. The first factor was the relaxed constraint on the required agreement of the output displacements of the Driving View and the Test View. Basically, because the last transition of a Test View to a Driving View was executed when the difference between their respective output displacements was within 40% of their lateral offset amounts, the IVLT technique relinquished control before the vehicle was fully centered in the destination lane. This means that control is relinquished when the vehicle is at a distance of up to 140% of the lateral offset distance away from the typical driving location in the destination lane. The second factor was the relatively loose closed loop vehicle control with respect to the Driving and Test View locations. Specifically, because the networks' output steering response to images created from Driving View and Test View were only compared to each other, and not to ground truth, transitions occurred before the vehicle had recentered itself in the Driving View image. For example, it was allowable for a transition to occur while driving on a straight road even if the Driving View and the Test View both responded with a moderate turn command. The reason that both networks produced this response was that the vehicle hadn't fully transitioned to the new driving location. i.e. The lane hadn't been recentered in the Driving View image. But because the output's were only compared to each other, the transition was allowed.

Figure 22 shows the average IRRE confidence values for the right and left lane ALVINN networks during right to left and left to right IVLT's. Like the confidence graphs associated with INS lane transitions, the IRRE confidence values associated with both the Driving and Test Views remained high throughout the transitions. The main feature that is visible in these graphs is the jump in confidence values, first visible for the Test View, and a short time later in the Driving View. This jump corresponds to the midway point of the IVLT, where the system changes which network is associated with the Test and Driving Views. For example, for a right to left lane transition, the right lane network is initially associated with the Driving View and the Test View. When the IVLT reaches its midway point, the left lane network first becomes associated with the Test View and then, if the constraints are satisfied, with the Driving View. This change corresponds to the jumps.

3.4.5 Dual View Lane Transition

Like the IVLT method, the Dual View Lane Transition method (DVLT) uses two virtual camera views to move the vehicle from the driving lane to the destination lane. While the IVLT method tracks only the original driving or destination lane to execute the lane transition, the DVLT method tracks them both. The idea behind the DVLT technique is to use a geometric road

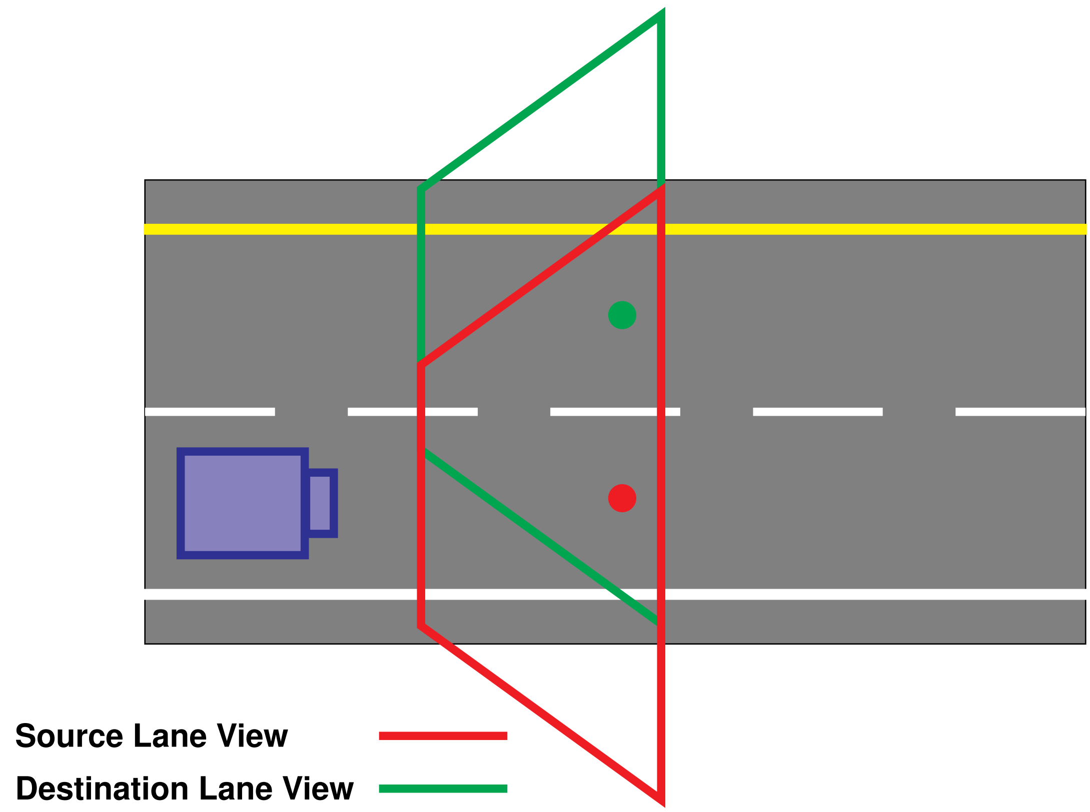

Figure 23 Initial virtual view placement for the DVLT technique.

model along with accurate knowledge of the location of the center of each lane to smoothly servo the vehicle between lanes.

This technique is a bottom up approach to lane transition. Although it can use input from high level modules to initiate the lane transition, the geometry of the situation is what drives this method. The system locates the center of both lanes and then moves the vehicle based on these locations. This is in contrast to the IVLT method where ALVINN's response to the movement of virtual camera views is what causes the lane transition to proceed.

Like IVLT, the DVLT method requires networks to be trained for the driving and destination lanes. When a DVLT is initiated, a second virtual camera view is created. The road model is used to laterally offset this view, called the Destination Lane View (DLV) so that it is centered over the destination lane. The network that was trained to drive in the destination lane is associated with

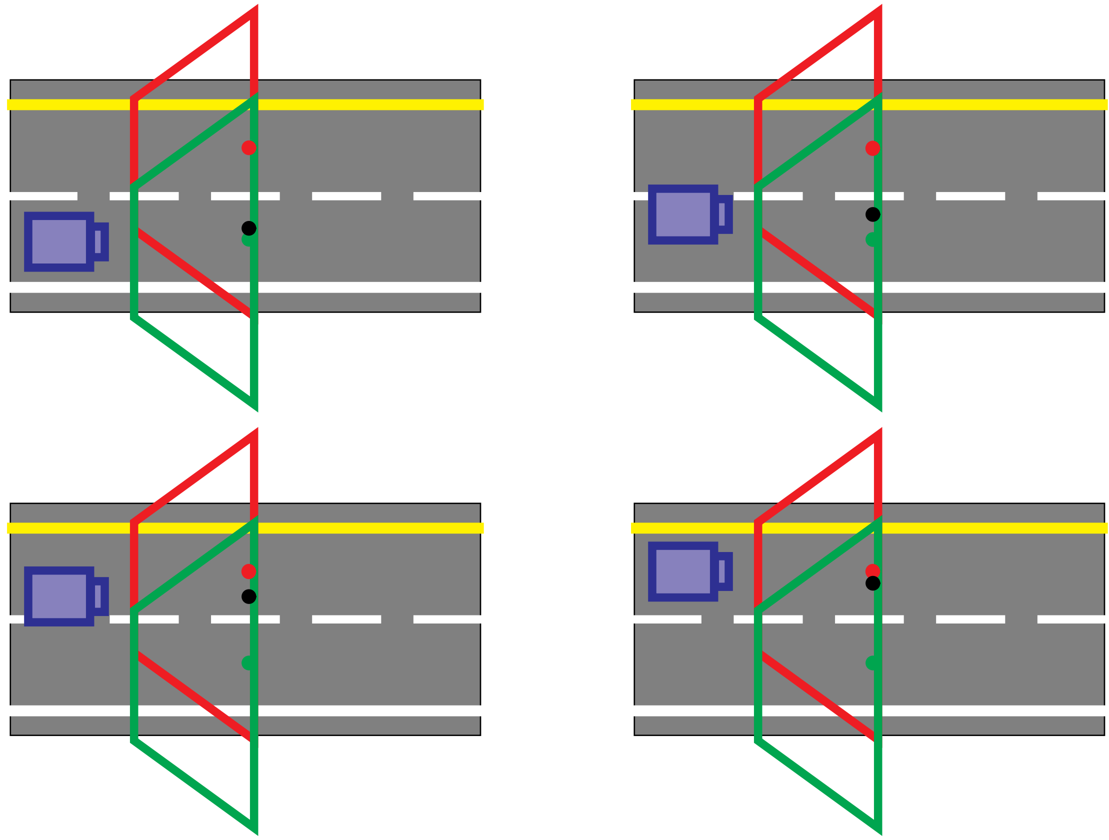

Figure 24 Movement of the vehicle and the views in the DVLT method.

the DLV. Contrary to the IVLT, the DLV is not treated as a Test View, but rather as a tracking device. The driving, or Source Lane View (SLV) is treated in the same manner. See Figure 23.

After the DLV has been created, images are generated from both the DLV and the SLV and are passed to their respective networks, which produce an output displacement along with an IRRE confidence value. Both output displacements are converted to vehicle relative points. These points, called Lane Center Points (LCP's), specify where ALVINN believes the center of each lane is located. Because each network has been trained using the same lookahead distance, and because the DLV is only offset laterally from SLV, the LCP's fall on a line which is perpendicular to the direction of travel.

At this point, the LCP's are used to calculate the Modified Lookahead Point (MLP). The MLP is the point that the vehicle will actually drive toward. The MLP is between the two LCP's, along the line connecting them. The distance between the MLP and either of the LCP's is related to the

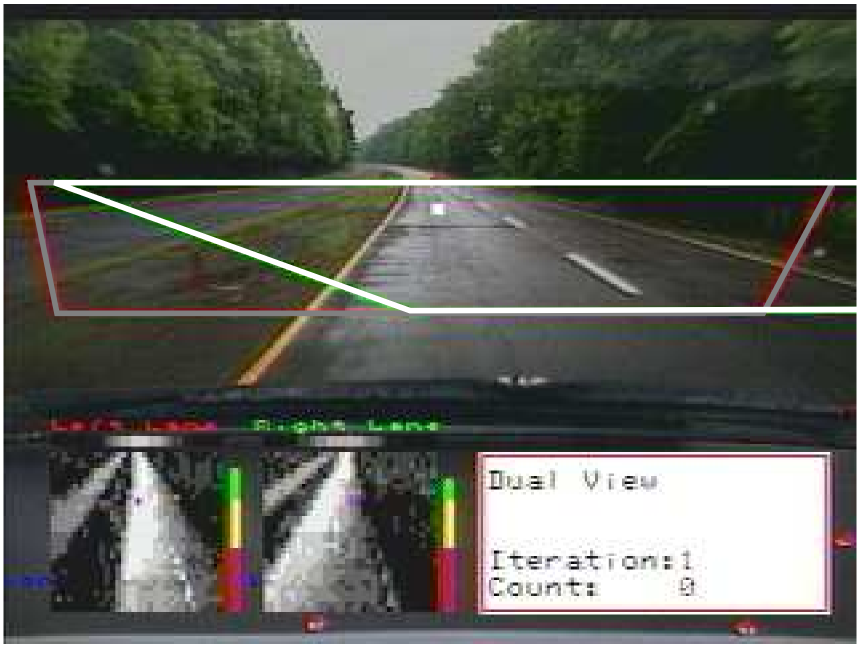

Left Lane Right Lane

Left Lane Right Lane

Figure 25 DVLT processing.

step in the DVLT process. For example, in DVLT experiments which use 16 iterations to transition between lanes, the first MLP would be 1/16 of the total distance (along the line connecting the LCP's) away from the driving lane LCP and 15/16 from the destination LCP. The second MLP would be 1/8 and 7/8, respectively. Because the vehicle is instructed to steer toward the MLP, the virtual views become misaligned with their respective lanes and must be updated. This process is similar to the one used for the IVLT method except that both views are changed. If the vehicle is moving to the left, both views are updated by moving them to the right. This process is continued until the vehicle has transitioned completely into the destination lane. See Figure 24. The black dot in each of the diagrams in this figure is the MLP. The combination of moving toward a MLP which is continually closer to the destination lane center and of updating the location of the virtual camera views results in a smooth, controlled transition.

During the lane transition, the IRRE confidence metric and the constant lane separation constraint are used to determine if the system is confident in its current driving ability. This is done by checking the IRRE confidence value associated with the images created by the views and by comparing the distance between the two LCP's. For the transition to continue, the IRRE confidence measure is required to be above a threshold value, typically 0.40, and the distance between the LCP's is required to be within 40% of the lane separation distance, specified by the lane model. Unlike the IVLT, this method does not check to see if the system will be able to drive at the next lane transition position.

Figure 25 shows two images taken at the beginning and end of a left to right DVLT. In the top image, the vehicle is still in the initial driving lane. The SLV footprint is visible as the darker trapezoid immediately in front of the vehicle. The preprocessed image associated with the SLV is shown in the lower left corner. The DLV footprint is also visible as the lighter trapezoid, offset to the right of the SLV footprint. The preprocessed image associated with it is shown to the right of the preprocessed image from the SLV. Just above each preprocessed image is ALVINN's driving response to that image. For each view, the driving response is almost the same, indicating that the constant separation distance road model and virtual camera imaging were accurate. Also shown in front of the vehicle is a grey dot which represents the MLP. Next to each preprocessed image is a bar graph representation of the IRRE confidence value associated with each. For both the SLV and DLV, the confidence is high - above 0.70.

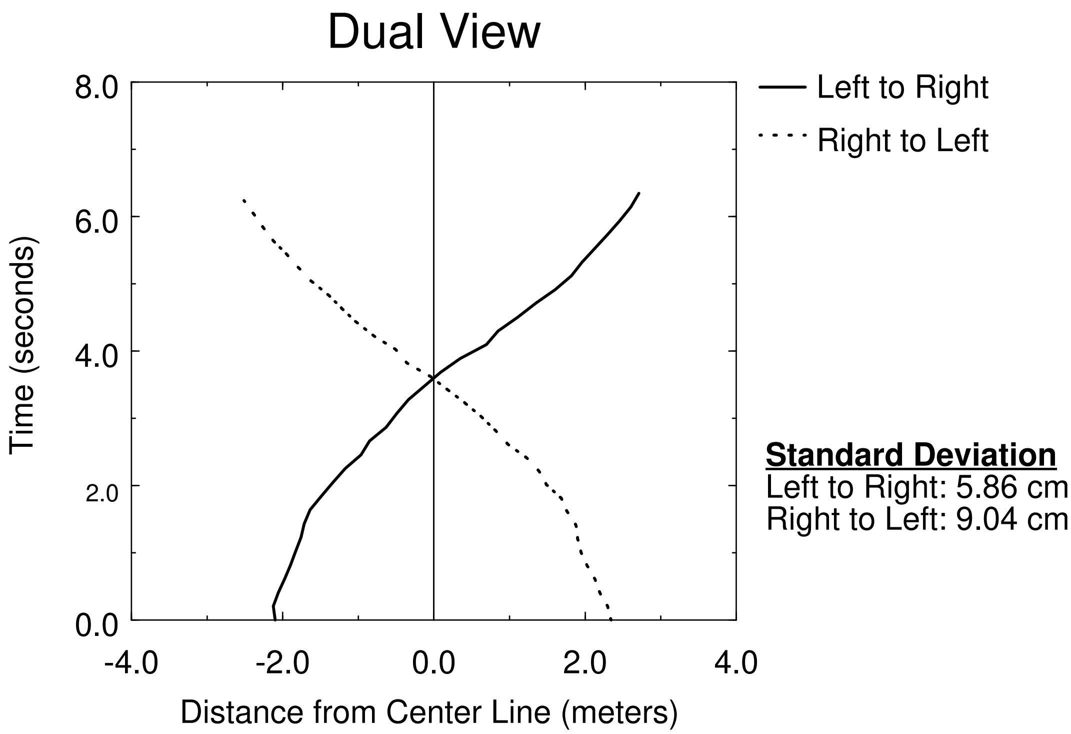

Figure 26 Average vehicle path during Dual View Lane Transitions.

The image on the bottom was taken as the vehicle approached the destination lane. In this image, the DLV footprint is almost centered in front of the vehicle while the original SLV is offset to the left. The preprocessed images are shown in the same location, along with ALVINN's driving responses and IRRE confidences values. The responses and confidence values are similar to those in the top image, even though the vehicle is now in the destination lane. The MLP is also shown.

All 42 DVLT's that were attempted to verify this lane transition method were successfully completed. Like the other methods, half were left lane to right lane transitions while the rest were right to left. The results of the experiments are shown in Figure 26. (The average distance for completion of an DVLT was 138 meters, while the average vehicle speed was about 21.9 meter/second.)

There are two characteristics to notice in this figure. First, the lane transitions are symmetric across the center line of the road as well as with respect to the time in the lane transition. Second, the DVLT method doesn't relinquish control until the vehicle is at the proper driving position in

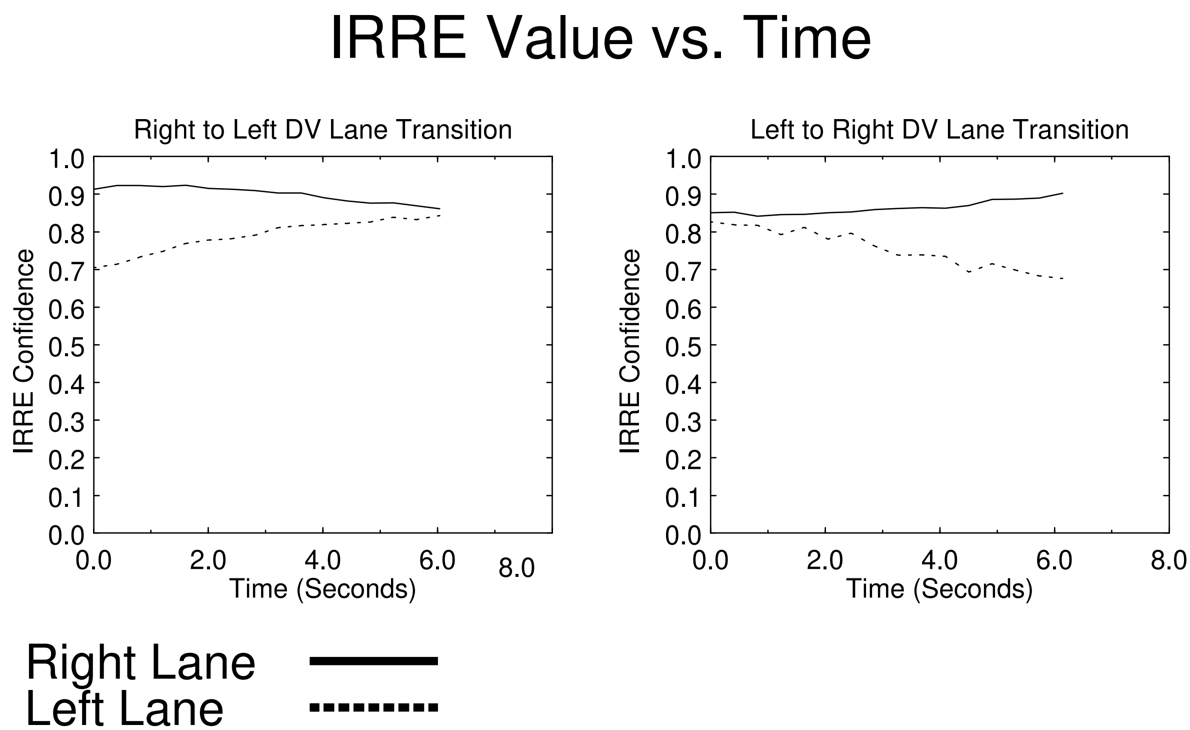

Figure 27 IRRE values during DVLT's.

the destination lane. Also notice that at the end of the transition the vehicle trajectory is still moving somewhat outward. A small amount of overshoot did occur and can be attributed to the existing lateral movement of the vehicle coupled with the moderately long response time that is a side effect of the large lookahead distance required for highway driving. But empirically, the DVLT method yielded the smoothest and most realistic transitions into the destination lane.

Figure 27 shows the average IRRE confidence values for the right and left lane ALVINN networks during right to left and left to right DVLT's. Similar to the graphs in Figure 18, the IRRE confidence values produced by an ALVINN network increase as the vehicle moves into the lane for which it was trained and decrease as the vehicle moves out of it. In both graphs though, the IRRE values remains well above the low confidence threshold. One cause of the increasing or decreasing confidence values is the extreme location of the SLV and DLV. Near the start and end of the transition, because each view remains centered over the proper lane, the actual camera viewing field may not overlay with the virtual camera viewing field. Although the missing virtual camera pixels are filled with the best actual camera pixel, the image is not quite consistent with what the network was trained with.

3.4.6 Lane Transition Experimental Parameters and Procedures

The same base virtual camera view was used to train all of the networks in each of the methods described. This means that each network "saw" the road or lane in the same manner. The virtual camera view used for these experiments was located 5.0 meters forward from the rear axle, 5.0 meters high, on the midline of the vehicle, and pitched downward from horizontal at an angle of 22 degrees. The actual camera was mounted inside the cab on the windshield, using the rear view mirror mounting post. It was located 2.11 meters forward from the rear axle, 1.36 meters above ground level, on the midline of the vehicle, and pitched downward at an angle of 10.67 degrees. Both the virtual and actual cameras were pointed straight ahead, i.e. they both had a yaw angle of 0 degrees. Also, all the networks were trained to drive with the same lookahead distance - 35 meters. Although using the same virtual view is not a requirement for any of the method to be described, training each network using the same lookahead distance is necessary.

The test site for all of the lane transition experiments presented here was the Pennsylvania Turnpike Commission's Safety Testing and Research (STAR) Facility in Breezewood, PA. The site is an abandoned, four lane section of the Pennsylvania Turnpike. For the experiments presented here, two lanes of the site were swept and repainted to look like a typical suburban interstate. Typical image of this road taken during a lane change maneuver are shown in Figure 28. All experiments were conducted on the Navlab 5 using the PANS platform. The vehicle speed was between 40 and 60 m.p.h.

To facilitate quantitative analysis of the lane transitions experiments, the live video signal used by ALVINN was recorded to videotape during the experiments. Overlayed on the video image was information regarding the type of lane transition and the step within each transition that was currently being executed. Network confidence and output displacement were recorded to disk along with the corresponding transition step. To determine the absolute position of the vehicle with respect to the center line of the road, the videotaped transitions were used. During each transition, the center line of the road was hand labeled.

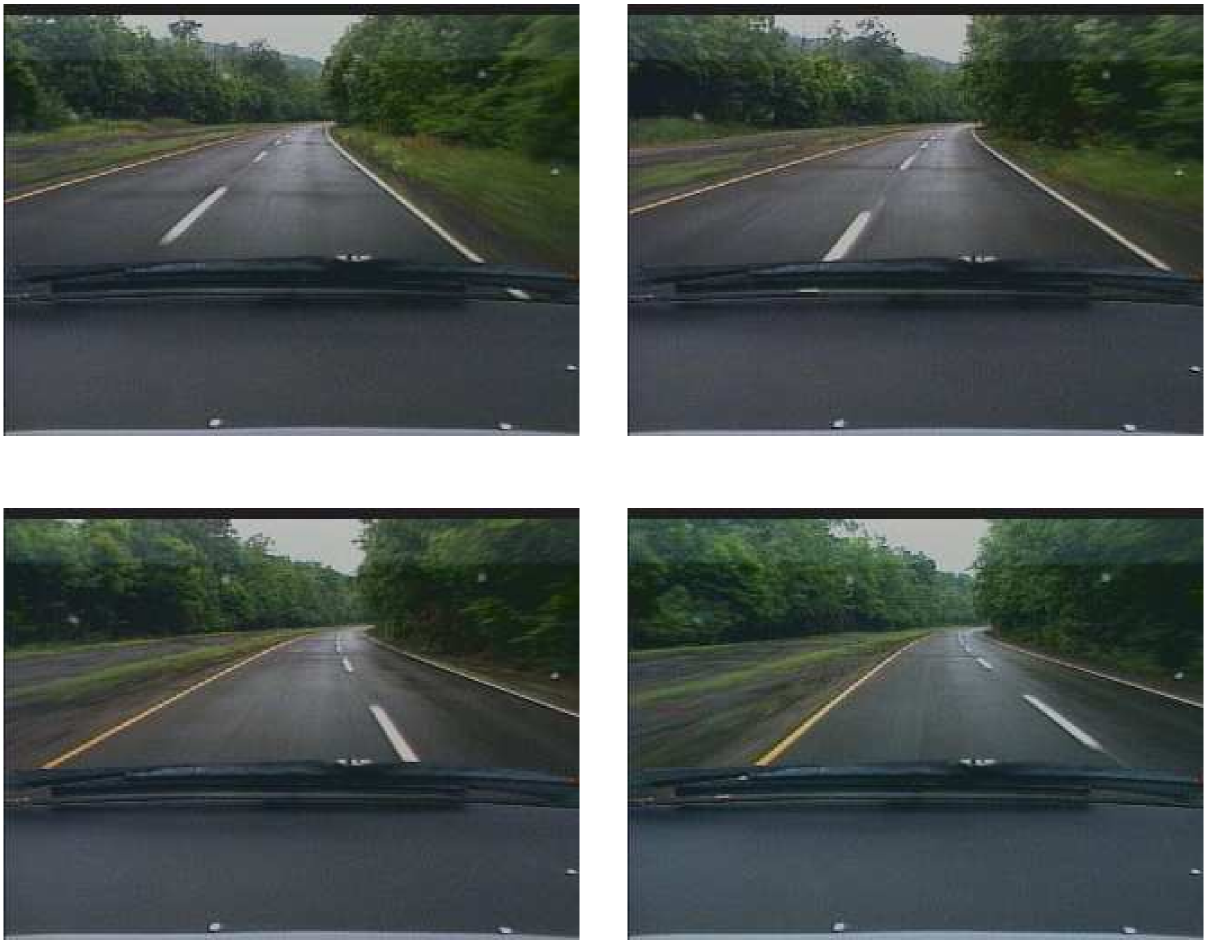

Figure 28 Pennsylvania Turnpike Commission STAR Site near Breezewood, PA.

3.5 Lateral Offset Driving

Although most systems are designed to keep the vehicle in the center of the driving lane, there are times when this behavior is not desirable. Obstacle avoidance is one such scenario. There are at least three reasons why moving the vehicle from the center of the driving lane when an obstacle is detected is important. The reasons relate to the capabilities of the obstacle detection system, the appearance of obstacle itself, and the capabilities of other cars on the roadway.

Many obstacle detection systems are currently being developed for use in the next generation passenger car. In these systems, which are mainly radar-based, the maximum detection distance of reasonably sized obstacles is very close to the minimum stopping distance for vehicles going the legal highway speed in the United States. If the systems do not detect the obstacle immediately and begin braking, the vehicle will not be able to avoid the obstacle by stopping.

The second reason why moving the vehicle from the center of the lane is important is that many obstacles will enter the driving lane at a distance closer than the minimum stopping distance of the vehicle. In this case, even immediate detection will not allow the vehicle to avoid a collision by stopping. In these cases, the only way to miss the obstacle is to swerve out of the current driving lane.

The final reason lane offset is important relates not to the capabilities of the vehicle equipped with the obstacle detection system, but rather to following vehicles which are not equipped with the systems. This type of "mixed-mode" traffic may be very common in the initial stages of a feasible Automated Highway System. In these situations, even if the detection system finds an upcoming obstacle in time to slow the vehicle to a stop, it may not be the safest maneuver because following vehicles may not have the same capability. They would not be able to stop and would either rear-end the stopped vehicle, or swerve out of the driving lane in order to avoid hitting it.

3.5.1 Swerve Experimental Results

Separate swerve systems do not need to be created to implement obstacle avoidance maneuvers. Active control of virtual camera views, similar to the techniques used for lane transition, can be used with current lane keeping systems to accomplish lane offset and swerve maneuvers.

To test this hypothesis, the IVLT technique presented earlier was modified into what is called the Swerve and Offset Driving Algorithm (SODA). The IVLT technique was selected as the base because it does not require two trained networks to initiate lateral movement of the vehicle, unlike the INS and DVLT techniques.

Instead of moving the vehicle completely into the destination lane like the IVLT technique does, SODA only offsets the vehicle from the center of the driving lane by a pre-specified amount, and then returns it to the normal driving position. This maneuver effectively allows the vehicle to swerve to miss an obstacle in the driving lane and is shown in Figure 29.

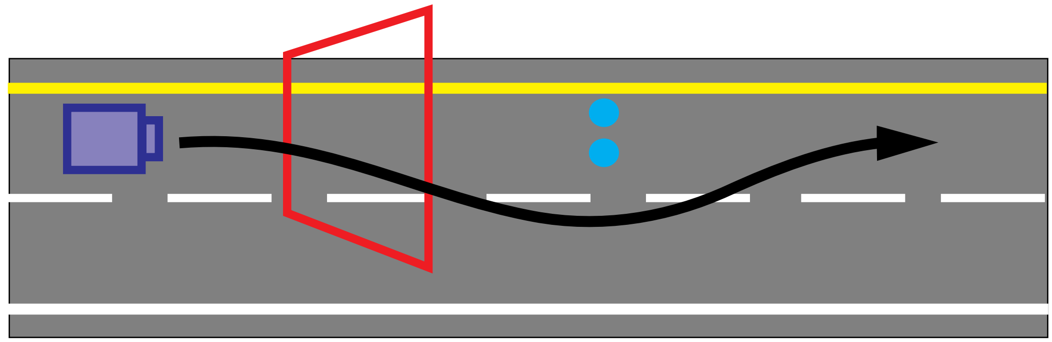

Figure 29 Vehicle movement while avoiding an obstacle.

SODA accomplishes the offset and return using a virtual camera control process, depicted in Figure 30, that is very similar to the one employed in the IVLT technique. The key differences are:

- Internal and external constraints are not enforced. This is done so that the maneuver takes place as quickly as possible. Although there is some lag between the virtual camera view location and the expected vehicle location due to the lack of coupling imposed by the constraints, the vehicle executes a smooth, stable maneuver.
- After the vehicle has reached the "apex" of the maneuver, it does not begin to return to the destination lane immediately, but rather drives with the specified offset for a predetermined distance. This "hold time" allows for some margin of error with respect to the actual location of the obstacle.

Approximately 10 left to right and 10 right to left swerves were attempted at the Breezewood test site. All were successfully completed. The obstacle was "detected" at distances between 40 and 60 meters by the test driver. After the test driver indicated an obstacle was present, the system moved the vehicle to an offset of 2.25 meters from the driving lane position using 8 virtual camera views. The system held the vehicle at this position for 0.5 seconds, and then returned it to its normal position in the driving lane. A condensed version of the maneuver, showing the virtual camera views and the vehicle's position relative to them, is shown in Figure 31.

3.5.2 Offset Driving Experimental Results

An additional experiment was conducted at the Breezewood site to test the feasibility of controlling the vehicle for an extended period while driving with a large offset from the lane which

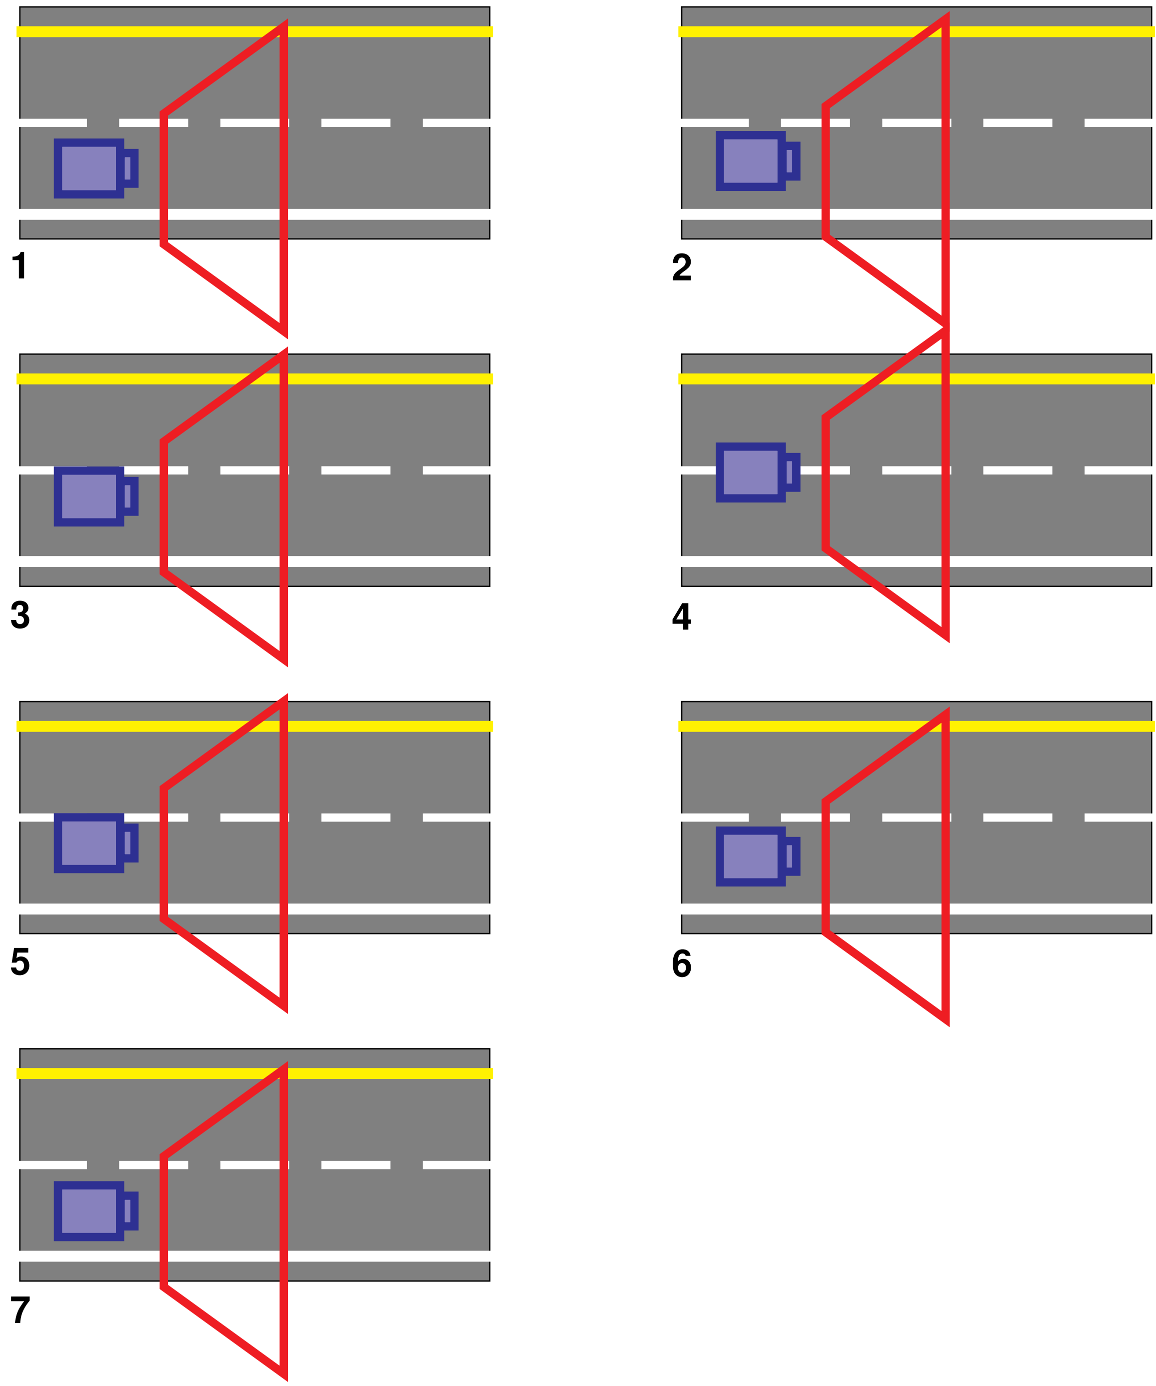

Figure 30 Vehicle and virtual camera locations during SODA maneuver.

the ALVINN network was originally trained. This type of behavior would be useful in circumstances where an emergency maneuver out of the driving lane and to a stop is necessary, such as when an accident is blocking the road ahead and the shoulder is the only escape route. The tech-

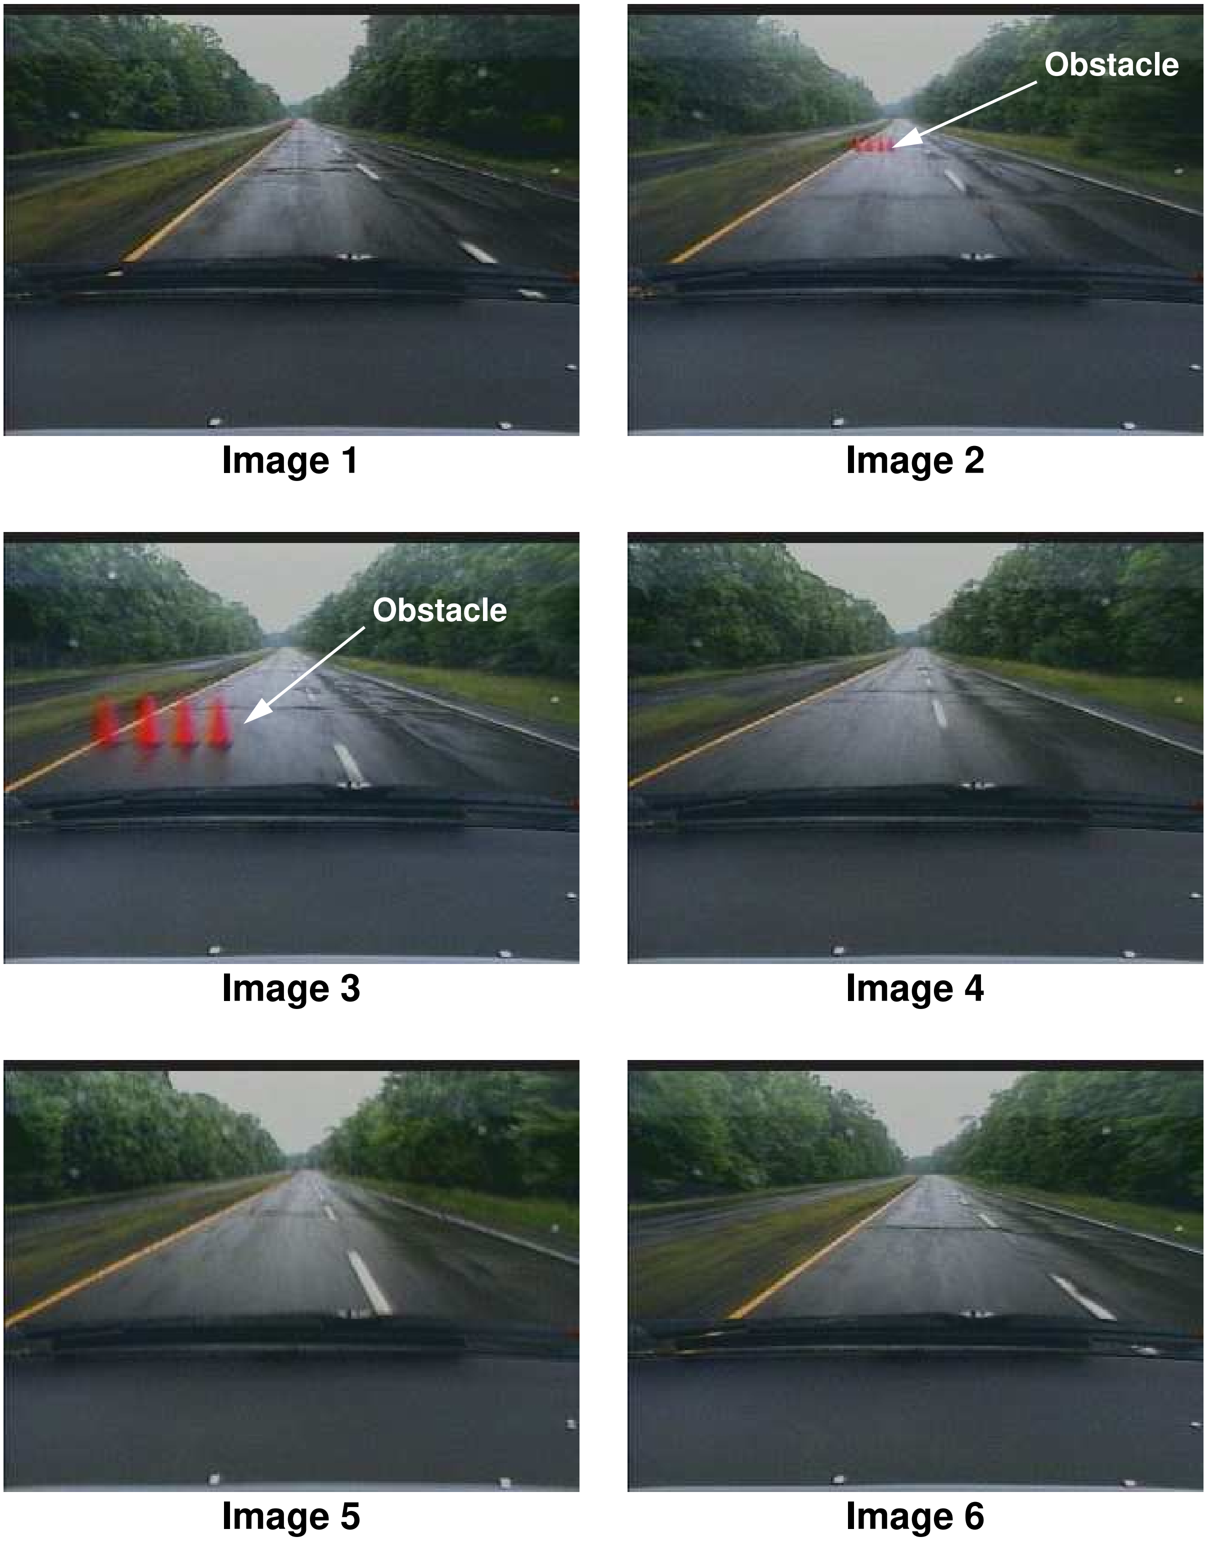

Figure 31 In-vehicle view of SODA maneuver.

nique could also be used in situations like road construction, where driving in the breakdown lane is necessary.

This maneuver was accomplished using the SODA algorithm, but by increasing the hold time to about 25 seconds. In this mode, the system was able to control the vehicle for over 580 meters (until the hold time had expired) and return it to the driving lane. Although performance was not as stable as when driving in the normal lane, it was well within acceptable limits, especially considering that this type of maneuver would be executed only as a last alternative.

3.6 Exit Ramp Detection and Traversal

A scenario which requires similar vehicle control schemes is exit ramp detection and traversal. This technique differs from simple lane transition because the exact location of the exit ramp is not usually known with enough precision to blindly move the vehicle onto it. The philosophy that has been adopted to accomplish this task is to:

1. Depend on other modules with higher level knowledge, like a map, to signal ALVINN when an exit ramp is approaching.
2. Use a virtual camera to preview the area of the scene where the exit ramp is expected.
3. Monitor the internal and external confidence measurements of the driving and exit lane networks which indicate the presence of an exit ramp.
4. Move the vehicle onto the exit ramp using lane transition techniques similar to those described earlier.

To accomplish this task using ALVINN, at least two trained networks are needed - one for the driving lane and another for the exit lane. (A third network may be needed to drive the vehicle down the exit lane once the vehicle has left the highway.) The road scenario is a two lane interstate highway where the exit lane is situated to the right of the right lane. ALVINN is driving the vehicle in the right lane.

Detection of the exit lane is the first sub-task which needs to be accomplished. Some distance before the exit lane begins, ALVINN receives a signal from a high level module informing it that the exit is approaching. Information about the exit lane's location relative to the vehicle, and type is passed to ALVINN. The information is used to create an appropriate virtual camera view, the

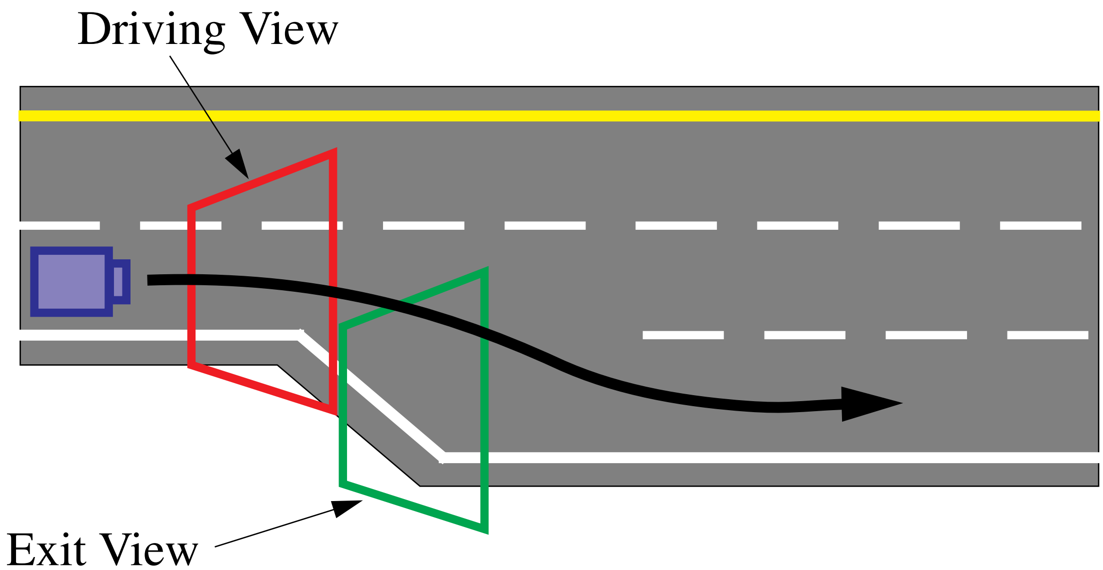

Figure 32 Exit ramp detection diagram.

Exit View, and associate the correct network with it. In addition to being laterally offset from the Driving view, the Exit View is also shifted forward by 10 meters. This is done so that the exit lane can be detected before the vehicle is adjacent to the lane and is important because in some situations the driving lane network learns to key off features, like the shoulder to road boundary, which change significantly when the exit lane appears. To maintain proper control in these situations it is important to detect the lane before it is imaged by the Driving View. A diagram of the roadway near the exit and virtual camera view locations is depicted in Figure 32.

While the system uses the Driving View to keep the vehicle in its lane, images are created from the Exit View and passed to the associated ALVINN network which creates an output steering displacement and IRRE confidence measure. By monitoring the confidence and output steering displacement of the exit lane network using the methods described in earlier sections of this chapter, the system is able to determine when the exit lane begins. The IRRE confidence of the exit lane network will become high and its output displacement will be the same as that of the driving view when the exit ramp is present. (The displacements will match because the driving and exit lanes are parallel and the Exit View is not rotated with respect to the road.)

Note that for this scenario, high level knowledge does not tell the system where or how to exit, but only what type of exit is approaching. Figure 33 shows the scene, virtual camera image, and network response when the exit lane is and is not present. The Driving View is outlined in black and the preprocessed image created from it along with its associated output displacement vector and bar graph representation of its IRRE measure is shown in the lower left corner of each image. The Exit View is outlined in white and the items associated with it are to the right of those for the Driving View. In the top image, where the exit lane is not present, both the internal and external constraints are not satisfied, namely the IRRE confidence value is low and the output displacement vector does not match that of the Driving View. In the lower image where the exit lane is present, both constraints are satisfied, which is indicated by the higher Exit View IRRE value and the nearly identical output displacements from both networks.

After the exit lane has been detected, the driving lane network is no longer used and the vehicle is controlled using only the output of the exit lane network. The Exit View is incrementally shifted both laterally and longitudinally toward its standard location in front of the vehicle. The lateral component of each shift results in an image in which it appears that the vehicle is offset to the left of its proper driving position in the exit lane. To recenter the vehicle, the network produces an output indicating that the vehicle should steer to the right, recentering it with respect to the image and moving it further into the exit lane. Like for IVLT's, the network confidence for each new virtual view is required to be above a threshold for two iterations before the view is shifted again. This process continues until the Exit View has reached its standard position and the vehicle has completely transitioned into the exit lane.

3.6.1 Exit Ramp Detection and Traversal Experimental Results

Of the 20 attempts at detecting and traversing exit lanes on a rural interstate highway, 19 were successful. The attempts were approximately evenly spread across three different ramps. The same driving and exit lane networks were used for every attempt, except for the single failure, in which a different exit lane network was used (and was the cause of the failure). Exit ramp detection occurred at approximately 55 m.p.h., while traversal was done at speeds between 35 and 50 m.p.h.

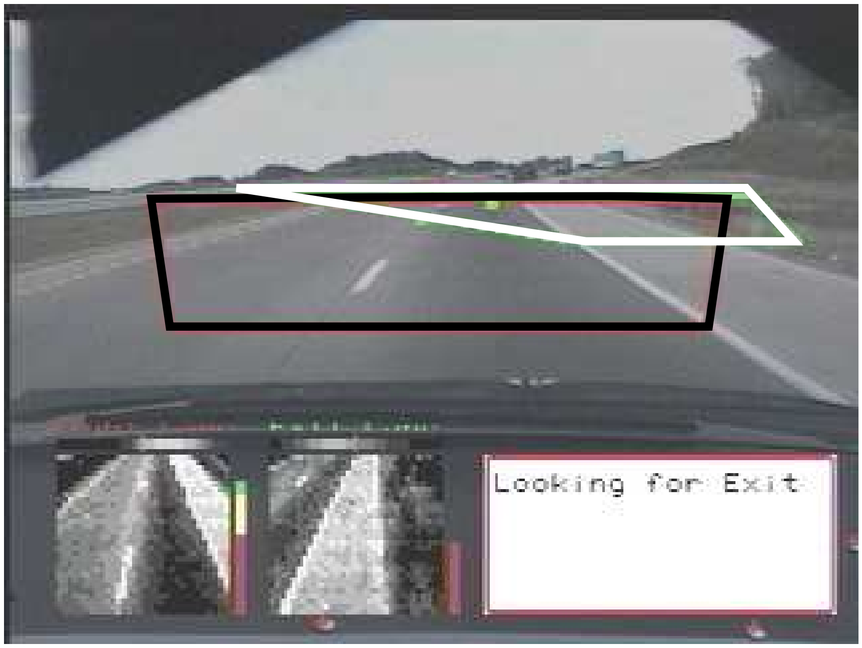

Driving Lane Exit Lane

Driving Lane Exit Lane

Figure 33 Exit ramp detection images.

Exit lane detection was robust and consistent, occurring immediately after the lane appeared centered in the Exit View. The system was instructed to begin looking for the exit lane between 100 and 800 meters before it actually occurred, and to begin transitioning onto it as soon as it was found. Finally, there were no false positives, in which the system incorrectly believed that the exit lane was present.

The system performance differed the most in the traversal part of the task. This was a result of the exit lane network and view that was used. This view was positioned so that the network would learn the most visible feature in the scene - the shoulder to offroad boundary. (The white lines on the concrete roadway were not a large enough feature and the exit lane was too lightly travelled for noticeable oil spots to be created.) The result of this choice of view meant that when the shoulder width changed, navigation performance would suffer. In practice, this lead to consistent overshoot during navigation of one of the three exits in which the shoulder was wider than that on which the network was trained.

3.6.2 Mainline Detection and Traversal

A task which is almost identical to exit ramp detection and traversal is mainline detection and traversal. This task refers to the act of moving along an entrance ramp, locating the right hand driving lane of the highway and transitioning onto it. It differs from exit ramp detection and traversal in that the view which is looking for the right hand lane is offset to the left of the diving view instead of the right and that after transition has occurred, the vehicle will be in a lane which has neighboring lane to the right and left.

Experimentation with this maneuver was conducted on a single entrance ramp using the same networks that were used for the exit ramp experiments. The exit lane network was used to drive the vehicle up the entrance ramp and the driving lane network was used to detect and travel onto the right lane of the interstate highway. The same algorithm used for exit ramp detection and traversal was used to find and move the vehicle into the right hand driving lane.

Figure 34 shows two images both before and after the right hand lane of the interstate highway was detected.

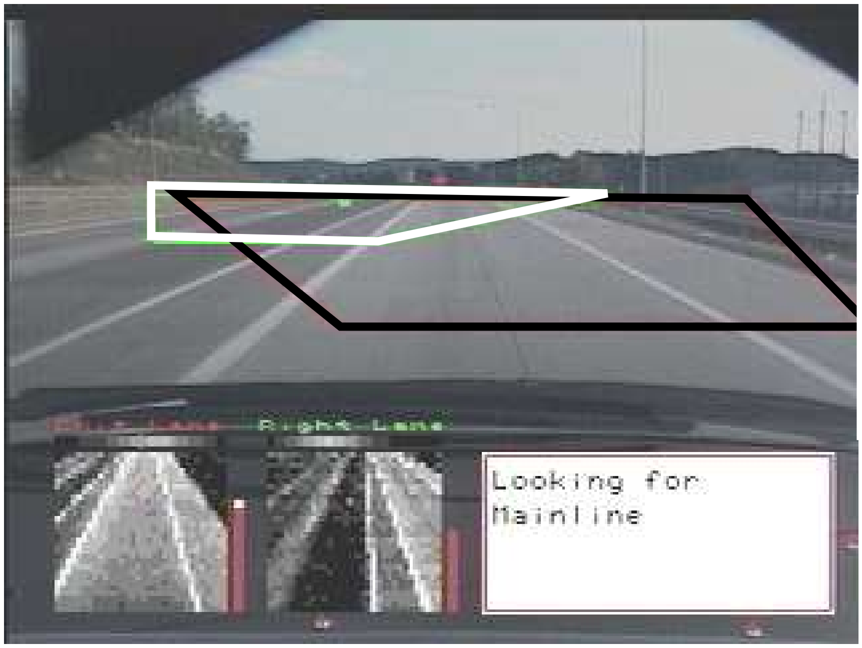

Driving Lane   Entrance Lane

Driving Lane   Entrance Lane

Figure 34 Entrance ramp detection images.

3.7 Lane Detection

In addition to being used to control the vehicle, virtual camera views can be used with ALVINN to signal the presence or absence of adjacent driving lanes. The lane transition and exit/entrance ramp scenarios described earlier in this chapter are two examples where this capability was used to determine if an action was appropriate. This idea is based on using the two metrics that are available which measure the reliability and confidence of ALVINN's output and was partially explained in earlier sections of this chapter.

To look for another driving lane, a network trained for the lane to be checked is needed and a virtual camera is placed so that it properly images the hypothesis road or lane location. An image is created from the virtual camera and is given to the network trained for that road or lane. The network produces an output displacement along with its associated IRRE confidence in this output. If the IRRE confidence is low, meaning that the image is not familiar to the network, two things are possible.

1. There is no lane present in the area that the virtual camera imaged.
2. There is a lane present, but the network used with the image created from the virtual camera view is not capable of driving in it.

In either case there is a mismatch between what was expected and what was present. Although, low confidence values are fairly indicative of the absence of another lane or incorrect network usage, high IRRE confidence values cannot be generally thought of as absolutely indicating the presence of another lane. An example of this phenomenon is shown in Figure 35. In this figure, a virtual camera view is being used to determine if another lane is present to the left of the driving lane. The top image shows the scene and the preprocessed driving and preview images created from the virtual camera views. In this case, the lane is present and as expected, the IRRE confidence value, shown to the right of preview image, is high. The bottom image of the figure shows a similar scene, except that now no lane is present. In this case, the network still responds with a high IRRE value. To overcome this inconsistency, the geometric constraints of the problem can be used. The constraints are:

1. The location of the virtual camera view with respect to the vehicle is known.
2. The lane separation distance is known and assumed to be constant.

3. The road curvature is known from the current driving view.

These constraints mean that the output displacement of the network associated with the lane that is being checked can be validated against the actual driving displacement. Also visible in Figure 35 are the network output steering vectors for the driving and preview networks. For both images, the output steering vector of the associated network is shown immediately above the preprocessed image. In the top image, the output steering vector associated with the preprocessed image that contains the road is almost identical to the output of the network which is controlling the vehicle. But in the bottom image, in which the preview virtual camera is not imaging the road, the output displacement is very different. Because the shoulder looks similar to the road, the network responds with a high confidence value, but by examining the output displacement vector of the preview view, it is possible to determine that no lane is actually present.

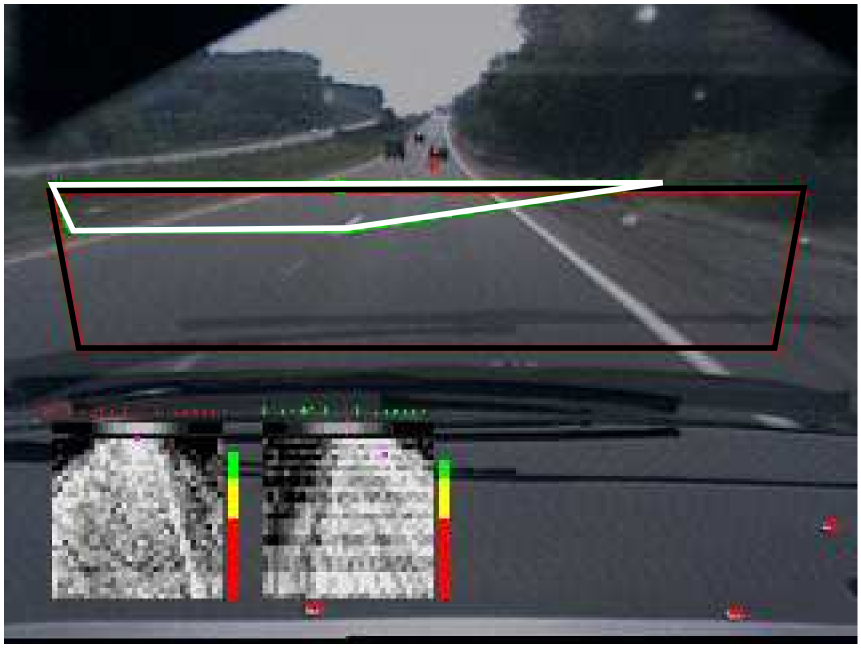

Driving Lane **Preview**

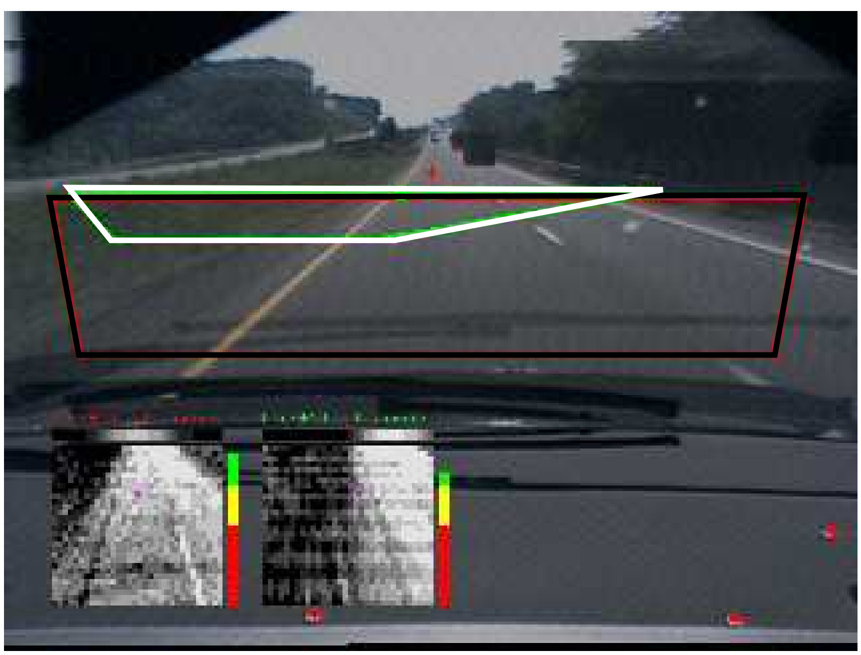

Driving Lane **Preview**

Figure 35 Lane detection scenario.